

Corporate · People Team (HR)

Performance Management Agent

CAI Score L3 · Operate — Bounded autonomy

▲ this agent

L1 Assist	L2 Draft	L3 Operate	L4 Decide	L5 Autonomous
--------------	-------------	---------------	--------------	------------------

TIER	PILLAR	AUTOMATION	HUMAN ROLE
L3 · Operate	Corporate	≈60% of routine effort	Owns exceptions and oversight

01 — Executive summary

Performance Management Agent is a CAI Score L3 (Operate) agent for corporate. Supports managers in performance management, helping structure reviews and keep feedback consistent across teams. It operates with bounded autonomy: a person owns exceptions and retains workflow oversight. Indicatively, it can redirect around 60% of the routine effort this work consumes.

02 — Problem & context

- Review quality swings wildly depending on the manager.
- Managers reinvent the structure of every review from scratch.

03 — Representative use cases


- Guides managers through the review structure, ensuring all required elements are addressed consistently.
- Surfaces relevant performance data and prior feedback for the manager before each review conversation.
- Reduces the time managers spend preparing for reviews by providing a structured starting point.
- Improves review quality consistency across the organisation by standardising the process, not the outcome.
- Produces the evidence to track review completion rates, feedback quality scores, and manager-consistency metrics.

04 — How to use it

- 1 Connect it in your tenant to the systems this work already touches (people team (hr)).
- 2 Confirm scope, the named owner, and the oversight the tier requires.
- 3 Detects the trigger in the workflow
- 4 Handles routine in-policy cases end to end
- 5 Escalates exceptions to the owner
- 6 Logs every action for audit
- 7 Review the audit log and KPI movement after each cycle; tune scope as you learn.

05 — Where it fits in the process


Sits inside the performance cycle — helping managers structure reviews and feedback so quality is consistent across the organisation.



Operating loop for a L3 agent – bounded autonomy.

06 — Risk assessment

#	RISK	LIKELIHOOD	IMPACT	MITIGATION / CONTROL
1	Incorrect or incomplete output	Medium	Medium	Routine cases only; exceptions and anything material are escalated to the named owner, with a full audit trail.
2	Over-reliance / reduced human scrutiny	Medium	Medium	Tier oversight keeps a human in the loop; KPIs monitor quality so reliance never outruns control.
3	Exposure of sensitive data	Medium	Medium	Runs inside your own tenant; Colleague AI processes none of your data, and access stays under your existing controls.
4	Behavioural drift over time	Medium	Low	Periodic review against KPIs and sampling detects drift; scope and prompts are version-controlled.
5	Out-of-policy or edge-case handling	Medium	Medium	Scope is bounded and documented; anything outside it is escalated, not improvised.



Residual likelihood × impact after the tier's built-in controls. Numbers map to the table above.

07 — Regulatory & compliance impact


FRAMEWORK	RELEVANCE	HOW THIS AGENT ALIGNS
EU AI Act	Operational use under human oversight; logging and risk-tiering apply.	Documented risk tier, human oversight and event logging per agent — designed to map to the obligations.
DORA	Relevant where the function is in scope for financial-sector operational resilience.	Actions are logged and attributable, supporting resilience and traceability requirements — designed to map to.
GDPR	Applies wherever personal data is in scope of the workflow.	Data stays in your tenant; Colleague AI processes none of it. Purpose limitation and minimisation are preserved.
ISO/IEC 42001	AI management-system controls: roles, monitoring, review.	The agent operates within a documented management system — defined tier, named roles, logging and review. Designed to meet; not certified.

“Designed to map to” / “designed to meet” describe alignment with the CAI Score framework, not external certification. Confirm obligations with your compliance function.

08 — Value & illustrative ROI

Supports managers in performance management, helping structure reviews and keep feedback consistent across teams.

Assumptions (illustrative — replace with your own)	Result
6 FTE engaged · 1840 productive h/FTE/yr · 60% addressable · €55/h blended	≈ 6,624 hours/yr redirected · ≈ €364,000/yr indicative value



Figures are an illustrative model to be validated against your own volumes and rates – not a guarantee.

KPIs to track

Improved management consistency.

09 — Recommendations

- Load the performance management framework and review structure into the agent before deployment.
- Managers own every assessment and decision — the agent structures and supports, it does not evaluate.
- Sensitive performance context stays with the manager and People Ops; the agent does not retain it.
- Review review-quality scores and completion rates after the first full cycle.

10 — Governance & method

This agent runs in your own Microsoft/Azure tenant; Colleague AI hosts only the control plane (score, policy, audit) and processes none of your data. A person owns exceptions and retains workflow oversight. Every action is time-stamped and attributable. The CAI Score classification is documented and open to challenge.

This dossier is an evaluation document. ROI and automation figures are illustrative models, not commitments. Compliance statements describe the CAI Score framework and are not a certification of your compliance with any law or standard. © Colleague AI — proprietary. CAI Score™.